arXiv:1512.05004v2 [cs.DL] 3 Feb 2016

Towards Cultural-Scale Models of Full Text
HTRC ACS Technical Report

Jaimie Murdock[1,2,*], Jiaan Zeng[1], and Colin Allen[2,3,4]

[1]School of Informatics and Computing, Indiana University, Bloomington, IN 47405, USA
[2]Program in Cognitive Science, Indiana University, Bloomington, IN 47405, USA
[3]Department of History and Philosophy of Science and Medicine, Indiana University, Bloomington, IN 47405, USA
[4]School of Humanities and Social Sciences, Xi'an Jiaotong University, Xi'an, China
[*]Corresponding author: jammurdo@indiana.edu

February 5, 2016

Abstract

This technical report consists of two components: an administrative report for the HathiTrust Research Center (HTRC) Advanced Collaborative Support (ACS) program and a research report on the variance of topic models trained over random samples of books in the Hathi Trust.

Keywords— digital libraries; topic modeling; topic alignment; random sampling; HathiTrust; Library of Congress Classification Outline (LCCO).

Administrative Report

"Taxonomizing the Texts: Towards Cultural-Scale Models of Full Text" was accepted to the HathiTrust Research Center (HTRC) Advanced Collaborative Support (ACS) program on Februrary 2, 2015. The first grant meeting was February 17. The ACS-supported data capsule deployment of the Topic Explorer was presented at the HTRC UnCamp 2015 on March 31 at University of Michigan. A tutorial on this deployment was presented at the Joint Conference on Digital Libraries (JCDL) 2015 on June 21 at the University of Tennessee - Knoxville [3]. The project preliminary research report was uploaded to the arXiv pre-print server on December 15, concluding this stage of the grant [2].

The grant had two original investigations: (1) Do random samples from within given LCSH categories yield significantly different topic models (TMs)? (2) Can Thomas Jeffersons taxonomic classification scheme for his library be predicted by TMs? After reviewer feedback, the grant was scaled back to focus on the first investigation. Duplicate removal notebooks were also cut from the grant. We modified the initial requirements from the Library of Congress Subject Headings to the Library of Congress Classification Outline. Documents can be multiply classified under LCSH, but only up to 5 categories, with inconsistent multiple tagging classes. The LCCO data was more readily available in the MARC-XML records held by HathiTrust and was more amenable to comparative analysis.

The Topic Explorer [1] was deployed to the data capsule on February 12. Continuing improvements were made to the data capsule deployment of the topic explorer throughout the summer for the UnCamp and JCDL with service releases on March 29, May 1, and June 12. The random sample generator and Library of Congress Classification Outline index was deployed on March 10. Subject area topic modeling was developed in Fall 2015, behind grant proposal work plan due to unanticipated challenges in topic model alignment and evaluation. Subject area topic modeling methods were completed on October 14, with additional improvements through December 5.

Community Impact

The Topic Explorer has become a core component of the HTRC Data Capsule. In addition to our research team's presentations at the HTRC UnCamp 2015, a tutorial at JCDL 2015, and at the HathiTrust User's Meeting, the work was presented by the larger HTRC community at the Humanities Intensive Learning and Teaching (HILT) conference in Indianapolis in July.

In addition, the work shows continuing promise. This semester an undergraduate research assistant with the InPhO lab has taken on scaling the experiments to all subject areas in the Library of Congress. We are submitting a proposal to the Institute of Museum and Library Services (IMLS) National Leadership Grants for Libraries Program for federal support to scale up this experiment to include a whole-library study. The grant pre-proposal is due on February 2nd, with announcements made in August 2016.

Also, This work will be submitted to the International Conference on Computational Social Science, to be held on June 23-26, 2016 at Northwestern University[1]. The InPhO Project group of Colin Allen and Jaimie Murdock are interested in continuing collaboration and funding opportunities with the HTRC to further advance this research.

Research Products

[1] Jaimie Murdock and Colin Allen. Visualization techniques for topic model checking. *Proceedings of 29th Association for the Advancement of Artificial Intelligence (AAAI-15)*, 2015.

[2] Jaimie Murdock, Jiaan Zeng, and Colin Allen. Towards Cultural-Scale Models of Full Text. Dec 2015. `http://arxiv.org/abs/1512.05004`.

[3] Jaimie Murdock, Jiaan Zeng, and Robert H McDonald. Topic Exploration with the HTRC Data Capsule for Non-Consumptive Research. In *JCDL '15 Proceedings of the 15th ACM/IEEE-CS Joint Conference on Digital Libraries*, Knoxville, Tennessee, USA, 2015. ACM Press.

[1] `http://www.kellogg.northwestern.edu/news-events/conference/ic2s2/2016.aspx`

Towards Cultural-Scale Models of Full Text
HTRC ACS Technical Report

Jaimie Murdock[1,2,*], Jiaan Zeng[1], and Colin Allen[2,3,4]

[1]School of Informatics and Computing, Indiana University, Bloomington, IN 47405, USA
[2]Program in Cognitive Science, Indiana University, Bloomington, IN 47405, USA
[3]Department of History and Philosophy of Science and Medicine, Indiana University, Bloomington, IN 47405, USA
[4]School of Humanities and Social Sciences, Xi'an Jiaotong University, Xi'an, China
[*]Corresponding author: jammurdo@indiana.edu

February 5, 2016

Abstract

Cultural-scale models of full text documents are prone to over-interpretation by researchers making unintentionally strong socio-linguistic claims without recognizing that even large digital libraries are merely samples of all the books ever produced. In this study, we test the sensitivity of the topic models to the sampling process by taking random samples of books in the Hathi Trust Digital Library within different Library of Congress Classification (LCC) areas. For each classification area, we train several topic models over the entire class with different random seeds, generating a set of spanning models. Then, we train topic models on random samples of books from the classification area, generating a set of sample models. Finally, we align topics from the sample models to the spanning models and measure the alignment distance and topic overlap. We find that sample models with a large sample size typically have an alignment distance that falls in the range of the alignment distance between spanning models. Unsurprisingly, as sample size increases, alignment distance decreases. We also find that the topic overlap increases as sample size increases. However, the decomposition of these measures by sample size differs by field and by number of topics. We speculate that these measures could be used to find classes which have a common "canon" discussed among all books in the area, as shown by high topic overlap and low alignment distance even in small sample sizes.

Research Introduction

Large-scale digital libraries, such as the Hathi Trust[1], give a window into a much greater quantity of textual data than ever before [6]. These data raise new challenges for analysis and interpretation. The constant, dynamic addition and revision of works in digital libraries mean that any study aiming to characterize the evolution of culture using large-scale digital libraries must have an awareness of the implications of corpus sampling. Failing to recognize that any large digital library is merely a sample of a larger set of books published within the culture can lead to unintentionally strong claims about socio-linguistics [10]. New methodologies also require careful consideration for humanistic implications [13].

One methodology with rapid uptake in the study of cultural evolution is probabilistic topic modeling [1]. Topic modeling has been used to characterize the evolution of literary diction [14], the evolution of literary studies [3], and to search large corpora for "the great unread" [12]. Moreover, topic modeling is an integral part of the Hathi Trust Research Center (HTRC)'s Data Capsule [9, 15].

Researchers need confidence in sampling methods used to construct topic models intended to represent very large portions of the HathiTrust collection. For example, topic modeling every book categorized under

[1]http://hathitrust.org/

the Library of Congress Classification Outline (LCCO)[2] as "Philosophy" (call numbers B1-5802) is impractical, as any library will be incomplete. However, if it can be shown that models built from different random samples are highly similar to one another, then the project of having a topic model that is sufficiently representative of the entire HT collection may become tractable.

Methods

LCCO Sampling

We implemented a random sampling web service that provides the following query interfaces:

- sampling(Category, number):
  It takes a category and the number of random samples as input, and returns a list of book ID which are randomly generated. For example, given category "DQ78-210" and 3, the web service returns "gri.ark:/13960/t50g6cm2v—uc1.31822038210555—uva.x030577307" where the book ID is separated by the pipe symbol;
- id(Category):
  It takes a category as input and returns a list of book ID which has all the books under such a category;
- idTotal(Category):
  It takes a category as input and returns the total number of books under such a category.

Figure 1 shows an example of the LCCO hierarchy stored in the HTRC Solr Index.

Corpus Download

Each subject area was downloaded from the HTRC on 19 October 2015 using the HTRC Data API. The selected areas can be found in Table 1.

LCCO	Subject Heading	# HT Vols
`nb`	Art Sculpture	801
`tg`	Bridge Engineering	799
`bj71-1185`	History of Ethics	898
`hd6050-6305`	Classes of Labor	1255
`tn600-799`	Metallurgy	938

Table 1: *LCCO Areas Sampled* — Areas in the Library of Congress Classification Outline (LCCO) and their representation in the HathiTrust Digital Library.

Topic Modeling

LDA topic modeling [2] represents the current state of the art for extracting meaningful data from digitized texts. We use the implementation of LDA embedded within the InPhO Topic Explorer [8], which uses collapsed Gibbs sampling for topic estimation [4]. All corpuses have the NLTK English stoplist removed. Additionally, all words occurring more than 50000 times and less than 15 times are removed from the corpus.

[2] `https://www.loc.gov/catdir/cpso/lcco/`

A *reference model* is trained on the whole subject area. Multiple other *spanning models* are trained on the whole subject area. For this preliminary study, we do not select the reference model from the spanning models, but research on model checking [1] and model selection [11] provide guidelines for further research. Finally, multiple *subcorpus models* are trained on different portions of the whole corpus, selected randomly.

Topic alignment

A topic alignment is a function that maps one topic in $\mathcal{M}_1$ to a topic in $\mathcal{M}_2$. For this analysis, $\mathcal{M}_2$ is always the *reference model*, while $\mathcal{M}_1$ is either a spanning or subcorpus model. The alignment function is not required to be bijective: multiple topics from $\mathcal{M}_1$ may map to the same topic in $\mathcal{M}_2$ (not injective) and not all topics in $\mathcal{M}_2$ are required to have an analog in $\mathcal{M}_1$ (not surjective).

Basic Alignment — The basic alignment simply computes the Jensen-Shannon distance (JSD) between each topic and all topics in $\mathcal{M}_2$. Each topic is simply matched to the closest topic, regardless of whether that topic in $\mathcal{M}_2$ is matched to another topic in $\mathcal{M}_1$.

Mathematically, this function is neither injective nor surjective, so an inverse alignment from M_2 to M_1 is not guaranteed to produce identical results.

Naïve Alignment — The naïve alignment also computes JSD between each topic and all topics in $\mathcal{M}_2$. However, if the closest match in $\mathcal{M}_2$ has already been selected by a topic in $\mathcal{M}_1$, the algorithm proceeds to the next closest topic until it finds one that has not yet been selected.

Mathematically, for equal numbers of topics in $\mathcal{M}_1$ and $\mathcal{M}_2$, the naïve alignment is bijective.

Results

In this preliminary study, we examined five areas of the LCCO. These five areas were selected to give examples in arts, humanities, sciences, and engineering and at different levels of granularity. For each subject area, three classes of models were trained at $k = \{20, 40, 60, 80\}$ topics:

1. LDA topic modeling was applied to all volumes in the subject area to generate the *reference model*.
2. Additional models were trained over all volumes to generate multiple *spanning models*.
3. Models were trained over random samples of the subject area to generate *subcorpus models*.

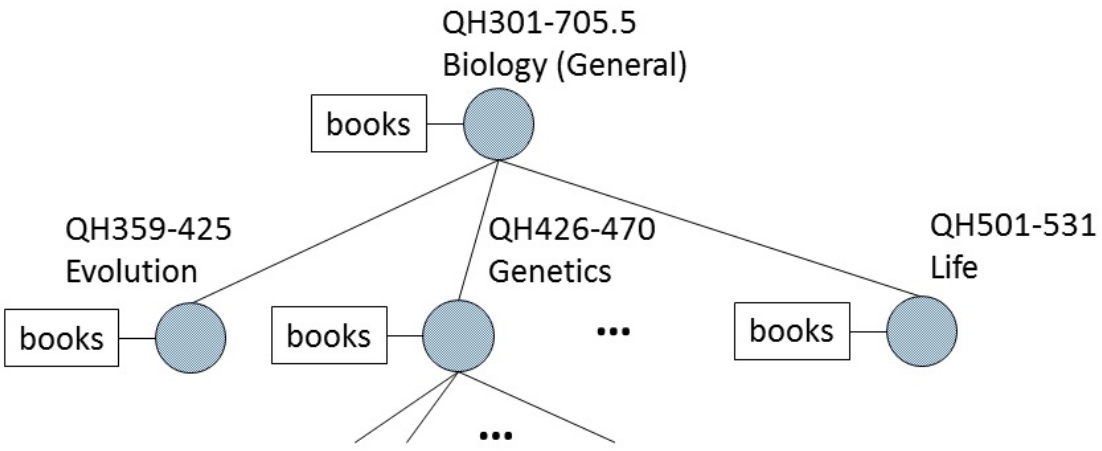

Figure 1: A sample tree representing the *QH301-705.5: Biology (General)* category.

A *basic alignment* was performed between each spanning or subcorpus model to the reference model. Each alignment's fitness score is calculated as the average Jensen-Shannon distance [5] between the word distribution of each topic pair. The results of the alignment are plotted against the sample size in Figure 2 for `nb`: Art Sculpture. Figure 3 shows the alignments for all five subject areas for $k = \{20, 40, 60, 80\}$.

Using the basic alignment, multiple topics in the spanning or subcorpus model may map to the same topic in the reference model and not all topics in the reference model need have an analog in the spanning or subcorpus model. We calculate the percentage of reference model topics captured by a basic alignment of the spanning or subcorpus model. Figure 4 shows topic capture for all five subject areas.

Discussion and Conclusions

In this preliminary study we wished to answer "Do random samples from within given LCCO categories yield significantly different TMs?" We find that the approach of principled subsampling can work, based on the existence of subcorpus models that have topic distances less than or equal to spanning models, relative to a reference model.

We highlight several phenomena for further investigation:

1. Topic distance decreases with sample size. (See Figures 2 and 3.)
2. Topic overlap increases with sample size. (See Figure 4.)
3. Different areas may have different characteristics, so a single sample size proportion may not be possible. (See Figures 3 and 4.)

In particular, the behavior of different areas may be tied to the "cognitive extent" of the discipline [7]. While we only used five subject areas, Figure 4 suggests differing levels of cognitive extent. We hypothesize that in fields such as Art Sculpture, History of Ethics, and Bridge Engineering, a canon has emerged that is common to the structure of many foundational textbooks in the discipline. For example, discussions of the History of Ethics are likely to discuss Aristotle, Kant, Hume, and Mill, regardless of the specific speciality. Discussions of Art Sculpture are very likely to discuss Michelangelo and Rodin, regardless of other influences. Discussions of bridge engineering are likely to start with a review of classical mechanics before diving into the equations necessary for a suspension bridge. If this hypothesis is true, there may be two types of classifications: those for emergent fields, and those for canonized fields. This is a new empirical question only made possible by corpora the size of the HathiTrust.

Acknowledgments

The work in this report was supported by a HTRC ACS Grant. We thank Miao Chen for her management of the ACS grants.We thank Justin Stamets for assistance with corpus curation.

References

[1] David M. Blei. Probabilistic topic models. *Communications of the ACM*, 55(4):77–84, April 2012.

[2] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.

[3] Andrew Goldstone and Ted Underwood. The Quiet Transformations of Literary Studies: What Thirteen Thousand Scholars Could Tell Us. *New Literary History*, 45(3):359–384, 2014.

[4] Thomas L Griffiths and Mark Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5228–5235, 2004.

[5] Jianhua Lin. Divergence Measures Based on the Shannon Entropy. *Information Theory, IEEE Transactions on*, 37(1):145–151, 1991.

[6] Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, The Google Books Team, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A Nowak, and Erez Lieberman Aiden. Quantitative Analysis of Culture Using Millions of Digitized Books. *Science*, 331(6014):176–182, Jan 2011.

[7] Staša Milojević. Quantifying the cognitive extent of science. *Journal of Informetrics*, 9(4):962–973, Oct 2015.

[8] Jaimie Murdock and Colin Allen. Visualization techniques for topic model checking. *Proceedings of 29th Association for the Advancement of Artificial Intelligence (AAAI-15)*, 2015.

[9] Jaimie Murdock, Jiaan Zeng, and Robert H McDonald. Topic Exploration with the HTRC Data Capsule for Non-Consumptive Research. In *JCDL '15 Proceedings of the 15th ACM/IEEE-CS Joint Conference on Digital Libraries*, Knoxville, Tennessee, USA, 2015. ACM Press.

[10] Eitan Adam Pechenick, Christopher M Danforth, and Peter Sheridan Dodds. Characterizing the Google Books Corpus: Strong Limits to Inferences of Socio-Cultural and Linguistic Evolution. *PLoS ONE*, 10(10):e0137041, 2015.

[11] Margaret Roberts, Brandon Stewart, and Dustin Tingley. Navigating the Local Modes of Big Data: The Case of Topic Models. In *Data Analytics in Social Science, Government, and Industry*. Cambridge University Press, New York, 2015.

[12] Timothy R. Tangherlini and Peter Leonard. Trawling in the Sea of the Great Unread: Sub-corpus topic modeling and Humanities research. *Poetics*, 41(6):725–749, 2013.

[13] Ted Underwood. Theorizing Research Practices We Forgot to Theorize Twenty Years Ago. *Representations*, 127(1):64–72, 2014.

[14] Ted Underwood and Jordan Sellers. The Emergence of Literary Diction. *Journal of Digital Humanities*, 1(2), 2012.

[15] Jiaan Zeng, Guangchen Ruan, Alexander Crowell, Atul Prakash, and Beth Plale. Cloud computing data capsules for non-consumptiveuse of texts. In *Proceedings of the 5th ACM Workshop on Scientific Cloud Computing*, ScienceCloud '14, pages 9–16. ACM, 2014.

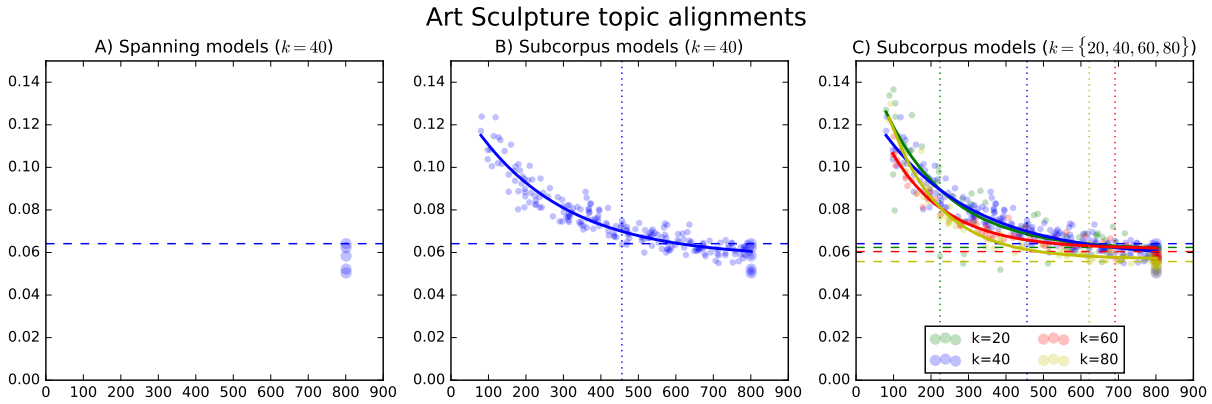

Figure 2: *Art Sculpture Topic Alignments*. Three cumulative plots demonstrate the approach over class `nb`. *A) Spanning models ($k = 40$)* — Five models trained against the entire corpus of 801 volumes are aligned to the reference model. The dashed line indicates the lowest-performing model. This model has an average Jensen-Shannon Distance of 0.064 bits per alignment pair. *B) Subcorpus models ($k = 40$)* — An additional 208 models are trained over samples of the whole corpus and aligned to the reference model. The solid line indicates an exponential decay regression fit to the data. The intercept between the worst performing model and the regression is at 618 volumes (77.1% of the total corpus size). The dotted line at 456 indicates the smallest subcorpus with performance exceeding the lowest-performing whole-corpus model (56.9%). *C) Subcorpus models ($k = \{20, 40, 60, 80\}$)* — The process in *B* is repeated for $k = \{20, 40, 60\}$.

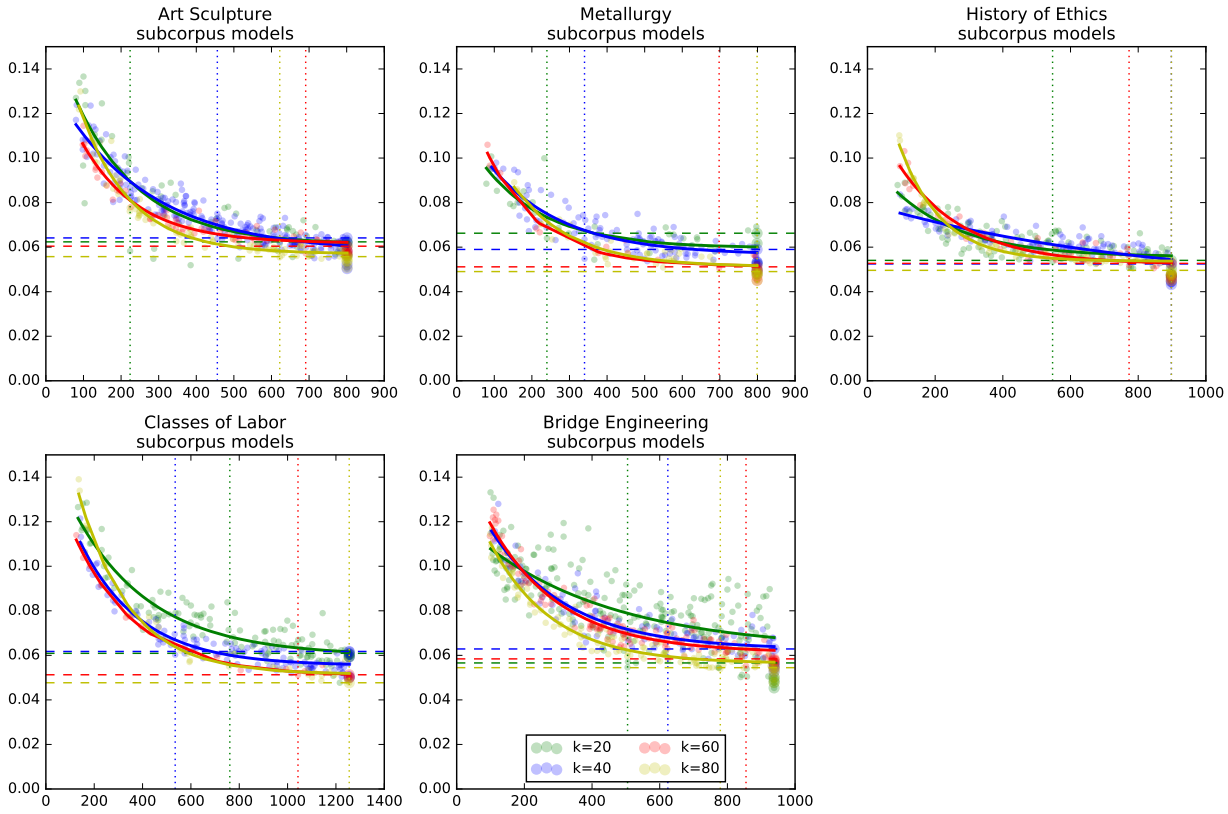

Figure 3: *Topic Alignments*. Subcorpus models ($k = \{20, 40, 60, 80\}$) for all five selected LCCO subject headings, replicating the analysis in Figure 2. In general, coarser models (i.e., those with lower number of topics) achieve the worst-whole-corpus-performance with a smaller subcorpus size than fine-grained models (i.e., those with higher number of topics).

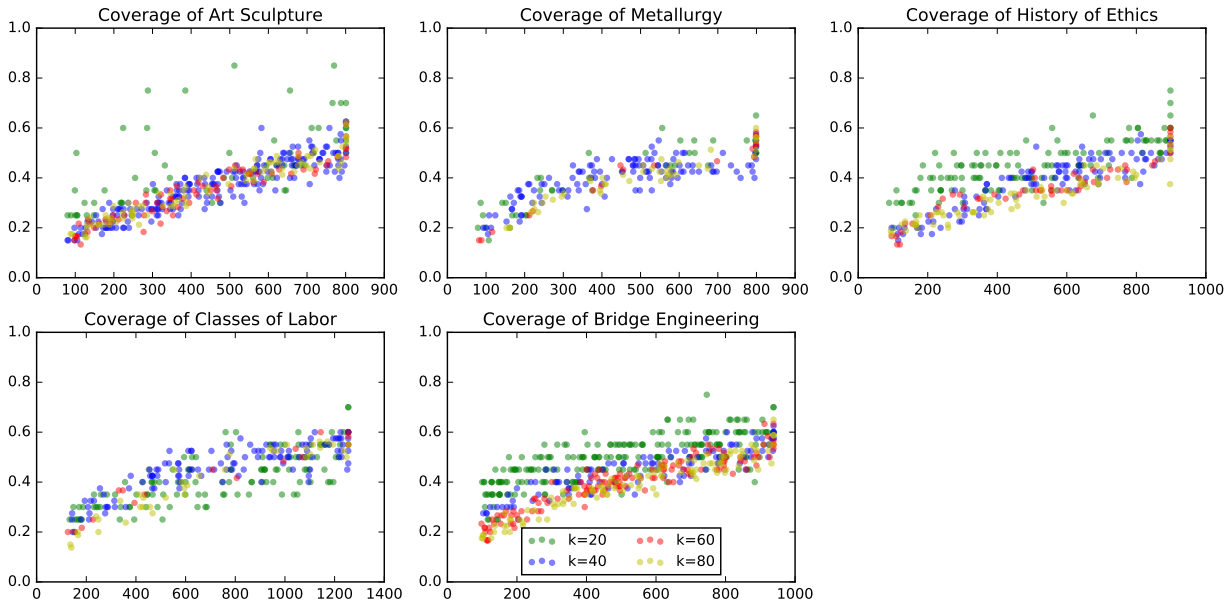

Figure 4: *Topic Capture* — Percentage of reference model topics captured by a basic alignment of the spanning or subcorpus model. Note that in Art Sculpture, History of Ethics, and Bridge Engineering the $k = 20$ models show higher topic alignment, in contrast to Metallurgy or Classes of Labor. See the Discussion for commentary on this observation.